



# Patch-Mix Contrastive Learning with Audio Spectrogram Transformer on Respiratory Sound Classification

Sangmin Bae<sup>1,5\*</sup>, June-Woo Kim<sup>2,5\*</sup>, Won-Yang Cho<sup>3,5</sup>, Hyerim Baek<sup>3,5</sup>, Soyoun Son<sup>5</sup>,  
Byungjo Lee<sup>4,5</sup>, Changwan Ha<sup>5</sup>, Kyongpil Tae<sup>5</sup>, Sungnyun Kim<sup>1,5†</sup>, Se-Young Yun<sup>1†</sup>

<sup>1</sup>KAIST AI <sup>2</sup>Department of AI, Kyungpook National University  
<sup>3</sup>SmartSound <sup>4</sup>Dongguk University <sup>5</sup>MODULABS

{bsmn0223, ksn4397, yunseyoung}@kaist.ac.kr kaen2891@knu.ac.kr

## Abstract

Respiratory sound contains crucial information for the early diagnosis of fatal lung diseases. Since the COVID-19 pandemic, there has been a growing interest in contact-free medical care based on electronic stethoscopes. To this end, cutting-edge deep learning models have been developed to diagnose lung diseases; however, it is still challenging due to the scarcity of medical data. In this study, we demonstrate that the pretrained model on large-scale visual and audio datasets can be generalized to the respiratory sound classification task. In addition, we introduce a straightforward Patch-Mix augmentation, which randomly mixes patches between different samples, with Audio Spectrogram Transformer (AST). We further propose a novel and effective Patch-Mix Contrastive Learning to distinguish the mixed representations in the latent space. Our method achieves state-of-the-art performance on the ICBHI dataset, outperforming the prior leading score by an improvement of 4.08%.

**Index Terms:** Respiratory Sound Classification, ICBHI, Audio Spectrogram Transformer, Patch-Mix Contrastive Learning

## 1. Introduction

Respiratory diseases, such as pneumonia or chronic obstructive pulmonary disease (COPD), impose a substantial burden on global health and well-being, constituting prevalent and common causes of death [1]. Furthermore, in the aftermath of the COVID-19 pandemic, there has been an increasing demand for accurate and expeditious diagnostic systems. The conventional approach for detecting lung diseases is auscultation, which utilizes a stethoscope to listen to the respiratory sounds of patients in person [2]. While hospitals can provide precise diagnoses, their effectiveness is significantly limited by the level of professional expertise and the subjectivity in interpretations [3, 4].

Recent developments in electronic stethoscopes provide recorded lung sounds, such as the ICBHI dataset [5], increasing the interest in contact-free medical care and opening up new possibilities for automated lung sound examination. Subsequently, numerous studies have developed methods for training neural networks that can achieve superior classification on the representative ICBHI dataset of stethoscopic sounds [6, 7, 8, 9]. Here, detection of abnormal respiratory sounds, such as *crackle* and *wheeze*, is critical for the accurate diagnosis of lung diseases [10]. In detail, crackles are characterized by discontinuous and non-tonal sounds that are associated with COPD, pneumonia, and lung fibrosis [11]. Wheezes, on the other hand, are high-pitched continuous sounds that are commonly detected in patients with asthma and COPD [10, 12].

\*equal contribution †corresponding authors

The code is available at: <https://github.com/raymin0223/patch-mix-contrastive-learning>




Figure 1: Overview of Patch-Mix Contrastive Learning.  $E_{ij}$  denotes the embedding of  $j$ -th patch from the  $i$ -th spectrogram, and  $P_j$  is  $j$ -th positional embedding. The class and distill tokens in the AST are omitted. Note that dashed lines indicate the stop-gradient operation.

Despite its significance, detecting abnormal respiratory sounds suffers from the scarcity of medical data. The collection of adequate data to train deep networks still remains a challenge, and the usage of collected data is often restricted due to patient privacy concerns. Besides, the annotation of lung sounds requires specialized medical expertise. Although previous literature has proposed various solutions, such as newly designed architectures [6, 9, 13], data augmentation techniques [6, 8, 14], or learning algorithms [8, 9, 15], these approaches have not yet achieved a satisfactory level of classification performance. In this work, we present an effective learning approach that utilizes mixing augmentations on the pretrained Transformer model.

So far, prior research on lung sound classification has proved that fine-tuning the pretrained conventional visual models is a straightforward and effective approach to address the shortage of data [8, 9, 14]. However, there is a lack of study on recent attention-based models such as Audio Spectrogram Transformer (AST) [16], which are being widely used in the audio domain [17, 18]. Moreover, most works have considered either only visual- or audio-pretrained models by ImageNet [19] and AudioSet [20], respectively. Since we convert lung sounds to Mel-spectrograms for the image-level classification, the image and audio pretrained model can be generalized to the respiratory sound classification task. Indeed, we demonstrate that

fine-tuning an AST model, pretrained on large-scale datasets of both domains, achieves state-of-the-art performance in respiratory sound classification, especially on the ICBHI dataset.

In order to alleviate overfitting on insufficient samples of the ICBHI dataset, several studies have also introduced data augmentation techniques such as label-aware concatenation [8] or Mixup [14]. Taking advantage of the AST model that patchifies the input image, we propose a straightforward Patch-Mix augmentation that randomly mixes patches (or tokens) between different instances (see Figure 1). However, respiratory sound classes contain a hierarchy such that when instances from the *normal* and *crackle* classes are evenly mixed, it is ambiguous to regard the mixture as equally distributed classes or just an abnormal crackle class. Therefore, the direct application of conventional Mixup [21] or CutMix [22] loss functions has limitations in the context of lung sound classification tasks.

To this end, we propose a novel Patch-Mix Contrastive Learning to distinguish mixed representations in the latent space, which enables the model to learn the extent to which the mixed representation resembles the original, irrespective of the label information [23, 24]. It is clearly different from the conventional contrastive loss [25, 26] in that we consider representations from mixed sources as positive pairs. In practice, our proposed methods achieved a new record 62.37% Score on the ICBHI dataset, improving by 4.08% compared to the previous best Score. We believe that our methods can be extended to various benchmarks and downstream tasks.

## 2. Related Works

### 2.1. Respiratory Sound Classification

Since the ICBHI 2017 challenge, numerous neural network-based methods have emerged for the lung sound classification task, particularly those that utilize or modify conventional residual block-based models [6, 7, 8, 9, 13, 14, 27]. Several studies have yielded promising results by utilizing pretrained models on ImageNet [8, 9, 14] or AudioSet [15], rather than training from scratch. Furthermore, to effectively address the issue of limited data, various learning protocols have been proposed, such as device-specific fine-tuning [8], task-specific co-tuning [9], or supervised contrastive learning [15]. Our study first leverages the innovative AST model, which deviates from conventional visual models, and attains state-of-the-art performance on the ICBHI lung sound dataset. Additionally, our proposed contrastive learning framework has shown synergy with the Transformer model that patchifies the audio spectrograms.

### 2.2. Mixed Representation Learning

Due to a shortage of lung sound data, previous research has proposed augmentation techniques such as label-aware concatenation [8] or Mixup in the respiratory audio [6, 14]. Mixing different samples, such as Mixup [21] or CutMix [22], is a popular augmentation technique that has been applied regardless of the modality. Recently, token-level CutMix methods combined with self-attention visual models have been proposed: TokenMix [28] assigns the target score based on the content-based activation maps of two images, and TokenMixup [29] proposes Mixup with guided attention on the Transformer architectures. In this study, we introduce a straightforward Patch-Mix technique that replaces randomly selected patches with those from other samples without any complex operation. Drawing inspiration from prior works [23, 24], we further formulate a mixed contrastive loss to distinguish the mixture in latent space, learning how much the mixed representation resembles the original.

## 3. Preliminaries

### 3.1. Dataset Description

We used the respiratory sound database of ICBHI 2017 challenge organized at Int. Conf. on Biomedical Health Informatics [5]. The ICBHI dataset comprises 6,898 respiratory cycles, which have a duration of approximately 5.5 hours and are officially split into a train set (60%) and a test set (40%). Note that the patients are divided into either the train or test set without overlap. Each breathing cycle is annotated as one of the four classes: *normal*, *crackle*, *wheeze*, and *both* (crackle and wheeze). The train set includes a total of 539 recordings taken from 79 patients, that consist of 1,215 cycles of crackles, 501 cycles of wheezes, 363 cycles of both crackles and wheezes, and 2,063 cycles of normal breathing. Similarly, the test set includes 381 recordings taken from 49 patients, with a total of 649 cycles of crackles, 385 cycles of wheezes, 143 cycles of both crackles and wheezes, and 1,579 cycles of normal breathing.

### 3.2. Preprocessing Details

Since all respiratory cycles in the ICBHI dataset have various sampling rates, we resampled all recordings to 16kHz. Besides, to ensure the same length of all samples, we set each duration to 8 seconds, corresponding to 798 frames. We truncated the samples for longer cycles while we duplicated the cycle with a fade in/out operation until the desired length [15]. Then, we converted the audio waveform to a sequence of 128-dimensional log Mel filterbank (Fbank) features, using the window and overlap size of 25ms and 10ms, respectively [16]. We also applied the standard normalization on the spectrograms with the mean and standard deviation of  $-4.27$  and  $4.57$ , respectively.

### 3.3. Training Details

In order to prevent overfitting, we utilized SpecAugment [30] with a maximum mask length of 160 frames for the time domain and 48 bins for the frequency domain. Here, we used the average values of each spectrogram for masking, and time-warping augmentation is excluded for simplicity. For the AST model, we used the Adam optimizer with a learning rate of  $5e-5$ , cosine scheduling, and a batch size of 8. We fine-tuned the pretrained AST model for only 50 epochs but trained for 200 epochs when starting from scratch. For other architectures, we used a learning rate of  $1e-3$ , a batch size of 128, and trained the model for 200 epochs. For the stable training, we applied the momentum update with the coefficient of 0.5 to all learnable parameters. Besides, except for when using mixing augmentation, we used a weighted cross-entropy loss, inversely proportional to the number of class samples [8, 15]. Note that we present the results of our experiments over five random runs.

### 3.4. Evaluation Metrics

Following the evaluation metrics of the ICBHI 2017 challenge, we evaluate the classification performance by Score, the average of specificity ( $S_p$ ) and sensitivity ( $S_e$ ). The specificity and sensitivity are defined as follows:

$$S_p = \frac{C_n}{N_n}, \quad S_e = \frac{C_c + C_w + C_b}{N_c + N_w + N_b} \quad (1)$$

where  $C_i$  and  $N_i$  are the number of correctly classified and total samples in class  $i \in \{normal, crackle, wheeze, both\}$ , respectively. In addition to the evaluation of the 4-class case, we also consider the 2-class scenario where crackle, wheeze, and both classes are grouped into a single class, labeled as *abnormal*. For simplicity, we utilize a trained classifier for the 4-class and subsequently compute the  $S_e$  for the 2-class version [9].

Table 1: Lung sound classification performance according to various architectures and pretraining datasets. All reported numbers are the results of simple fine-tuning on the ICBHI dataset. IN and AS indicate ImageNet [19] and AudioSet [20], respectively. † denotes the experimental results of our implementation. **Best** and second best results.

architecture	pretrain	$S_p$ (%)	$S_e$ (%)	Score (%)
EfficientNet†	-	70.92±3.61	28.93±3.05	49.93±0.59
ResNet18†	-	70.68±3.28	30.14±3.98	50.41±0.60
bi-ResNet [13]	-	69.20	31.12	50.16
ResNet-NL [6]	-	63.20	41.32	52.26
ResNet-Att [27]	-	71.44	36.36	53.90
bi-ResNet-Att [31]	-	67.13	<b>46.38</b>	56.76
CNN6 [15]	-	76.72	31.12	53.92
EfficientNet†	IN	<b>80.07</b> ±4.74	33.05±4.14	56.56±0.79
ResNet18†	IN	76.63±1.87	31.06±2.45	53.84±1.26
ResNeSt [14]	IN	67.80	37.10	52.20
ResNet34 [8]	IN	71.40	39.00	55.20
ResNet50 [9]	IN	76.33	37.37	56.85
CNN6 [15]	AS	70.09	40.39	55.24
AST†	-	72.61±14.90	30.58±7.97	49.60±1.75
AST†	IN	78.69±0.60	38.78±0.65	<b>58.73</b> ±0.22
AST†	IN+AS	77.14±3.35	41.97±2.21	<b>59.55</b> ±0.88

## 4. Methodology

### 4.1. Audio Spectrogram Transformer

The respiratory sound classification task is a challenging problem due to the limited number of samples available to train the large networks. As a result, prior works have utilized pre-trained conventional visual models on either ImageNet (IN) or AudioSet (AS) to overcome the scarcity of training samples. AudioSet contains over 2 million audio clips representing various sound types, even including respiratory sounds like breathing, cough, sneeze, or sniff [20]. In fact, Moummad *et al.* [15] achieved comparable results to the best Score using the AudioSet pretrained CNN6 model.

Accordingly, pretraining on both visual and audio domains would be a powerful technique for improving the performance of respiratory sounds. The Audio Spectrogram Transformer (AST) [16], a purely attention-based model, is sequentially pre-trained on ImageNet and AudioSet, demonstrating significant improvements in various audio downstream tasks. The self-attention mechanism also enables it to better capture long-range dependency in both time and frequency domains. In contrast to previous works that solely rely on convolution models, we adopted the *base* AST model pretrained on both ImageNet and AudioSet for the lung sound classification task.

Table 1 summarizes the classification Score for the 4-class evaluation across different architectures and datasets. In order to confirm the influence of the architecture and pretraining datasets, we fine-tuned the models without additional learning techniques. Note that we compared only the reported values in previous studies that used the official split. Although training the AST model from scratch is challenging, there was a significant improvement as it was more pretrained on large datasets, resulting in state-of-the-art performance with ImageNet and AudioSet pretraining. This suggests that pretraining on datasets from both domains indeed enhances the generalization of the AST for the lung sound classification task.

### 4.2. Patch-level CutMix on AST Architecture

We have also explored the data augmentation techniques to regularize models with augmented samples. The Mixup [21]

Table 2: Patch-Mix contrastive learning performance based on three factors: mix augmentations, negative pair types, and target representations.  $y$  and  $stop$  indicate the label and stop-gradient operation, respectively. **Best** and second best results. We highlight the final implementation choices.

method	mix aug.	neg.	target	$S_p$ (%)	$S_e$ (%)	Score (%)
CE	-	-	-	77.14±3.35	41.97±2.21	59.55±0.88
CE	Patch	-	-	76.87±5.05	42.04±3.98	59.46±0.78
CE	T-Patch	-	-	79.04±4.88	40.55±5.14	59.79±1.20
MixCL	Patch	ALL	stop(z)	<b>81.66</b> ±3.83	43.07±2.80	<b>62.37</b> ±0.61
MixCL	T-Patch	ALL	stop(z)	78.22±7.96	<b>44.79</b> ±6.17	61.51±1.15
MixCL	Patch	Diff. $y$	stop(z)	80.24±4.60	40.93±4.44	60.59±0.50
MixCL	Patch	ALL	stop(h(z))	78.07±2.41	<u>43.54</u> ±1.75	60.81±0.41
MixCL	Patch	ALL	$z$	<b>82.17</b> ±5.52	40.87±4.02	<u>61.52</u> ±1.37
MixCL	Patch	ALL	$h(z)$	79.89±4.08	42.06±4.75	60.97±0.55




Figure 2: Ablation study on hyperparameters of Patch-Mix contrastive loss. We searched hyperparameters in a default setting: Patch-Mix, ALL negative pair, and stop(z). We highlight the final chosen hyperparameters by gray color.

and Cutmix [22] methods are popular plug-and-play techniques that can be applied to any modality. Since Transformer-based models split inputs into a sequence of patches (or tokens), recent works have utilized mixing techniques at the patch-level [28, 29]. To build on this concept, we propose an augmentation technique called Patch-Mix that simply mixes patch embeddings with those from randomly chosen instances, as illustrated in Figure 1. Our approach does not rely on complex operations such as neural activation maps from a pretrained teacher model [28] or guided attention [29].

Initially, we adopted the same loss form as Mixup or CutMix, generating the label of cross-entropy loss based on the mixing ratio. Here, we experimented with two augmentation types of Patch-Mix: random selection of masks (“Patch” in Table 2) and masking among the patch groups belonging to the same time domain to mimic Cutmix in raw audio signals (“T-Patch”). However, both augmentations did not show the improvement compared to simple fine-tuning results. We hypothesize that the reason is the label hierarchy in respiratory sounds. Linear combinations of ground truth labels may be counter-intuitive in that patients should be diagnosed to *abnormal* with even small amount of crackle or wheeze sounds. Therefore, the conventional mixed loss function using  $\lambda$ -interpolated labels may not be appropriate for respiratory sound datasets.

### 4.3. Patch-Mix Contrastive Learning

In order to obtain a suitable loss function, we propose a Patch-Mix contrastive loss that enables the effective learning by distinguishing the mixed information within the latent space:

$$\mathcal{L}_{CL} = -\frac{1}{|I|} \sum_{i \in I} \left[ \left( \lambda \cdot (h(\tilde{z}_i)^\top z_i / \tau) + (1 - \lambda) \cdot (h(\tilde{z}_i)^\top z_m / \tau) \right) - \log \sum_{j \in I} \exp(h(\tilde{z}_i)^\top z_j / \tau) \right] \quad (2)$$

Table 3: Overall comparison on ICBHI lung sound classification task. We compared to previous works that use the official split of the ICBHI dataset (i.e., 60-40% split for the train and test set). \* denotes the previous state-of-the-art Score. **Best** and second best results.

	method	architecture	pretrain	venue	$S_p$ (%)	$S_e$ (%)	Score (%)	
4-class eval.	LungBRN [13]	bi-ResNet	-	<i>BioCAS'19</i>	69.20	31.12	50.16	
	SE+SA [7]	ResNet18	-	<i>INTERSPEECH'20</i>	81.25	17.84	49.55	
	LungRN+NL [6]	ResNet-NL	-	<i>INTERSPEECH'20</i>	63.20	41.32	52.26	
	CNN-MoE [32]	C-DNN	-	<i>JBHI'21</i>	72.40	21.50	47.00	
	LungAttn [27]	ResNet-Att	-	<i>IPEM'21</i>	71.44	36.36	53.90	
	ARSC-Net [31]	bi-ResNet-Att	-	<i>BIBM'21</i>	67.13	<b>46.38</b>	56.76	
	RespireNet [8]	ResNet34	IN	<i>EMBC'21</i>	71.40	39.00	55.20	
	RespireNet [8] (CBA+BRC+FT)	ResNet34	IN	<i>EMBC'21</i>	72.30	40.10	56.20	
	Ren <i>et al.</i> [33]	CNN8-Pt	-	<i>ICASSP'22</i>	72.96	27.78	50.37	
	Chang <i>et al.</i> [34]	CNN8-dilated	-	<i>INTERSPEECH'22</i>	69.92	35.85	52.89	
	Chang <i>et al.</i> [34]	ResNet-dilated	-	<i>INTERSPEECH'22</i>	50.22	51.83	51.02	
	Wang <i>et al.</i> [14] (Splice)	ResNeSt	IN	<i>ICASSP'22</i>	70.40	40.20	55.30	
	Late-Fusion [35]	Inc-03 + VGG14	IN	<i>EMBC'22</i>	<b>85.60</b>	30.00	57.30	
	Nguyen <i>et al.</i> [9] (StochNorm)	ResNet50	IN	<i>TBME'22</i>	78.86	36.40	57.63	
	Nguyen <i>et al.</i> [9] (CoTuning)	ResNet50	IN	<i>TBME'22</i>	79.34	37.24	58.29*	
	Moummad <i>et al.</i> [15]	CNN6	AS	<i>arXiv'22</i>	70.09	40.39	55.24	
	Moummad <i>et al.</i> [15] (SCL)	CNN6	AS	<i>arXiv'22</i>	75.95	39.15	57.55	
		<b>Fine-tuning [ours]</b>	AST	IN + AS	-	77.14 $\pm$ 3.35	41.97 $\pm$ 2.21	59.55 $\pm$ 0.88
	<b>Patch-Mix CL [ours]</b>	AST	IN + AS	-	81.66 $\pm$ 3.83	43.07 $\pm$ 2.80	<b>62.37</b> $\pm$ 0.61	
2-class eval.	CNN-MoE [32]	C-DNN	-	<i>JBHI'21</i>	72.40	37.50	54.10	
	Nguyen <i>et al.</i> [9] (StochNorm)	ResNet50	IN	<i>TBME'22</i>	78.86	49.79	64.32	
	Nguyen <i>et al.</i> [9] (CoTuning)	ResNet50	IN	<i>TBME'22</i>	79.34	50.14	64.74*	
		<b>Fine-tuning [ours]</b>	AST	IN + AS	-	77.14 $\pm$ 3.35	<b>56.40</b> $\pm$ 1.73	66.77 $\pm$ 1.15
		<b>Patch-Mix CL [ours]</b>	AST	IN + AS	-	81.66 $\pm$ 3.83	55.77 $\pm$ 2.83	<b>68.71</b> $\pm$ 0.59

where  $z_i$  is the encoder output of patch embedding  $E_i$ , and  $\tilde{z}_i$  is from mixed embedding with a ratio of  $\lambda$  between  $E_i$  and  $E_m$ .  $h$  is the projector that consists of two MLP layers with ReLU and BN layers. Note that all representation vectors, such as  $h(\tilde{z}_i)$  or  $z_i$ , are normalized before the dot product.  $I$  is the index set,  $\tau$  is the temperature to control the sharpness of cosine similarity, and  $\lambda$  is drawn from Beta distribution with parameter  $\beta$ . Here, we formulate a final loss function as follows:  $\mathcal{L}_{CE} + \alpha\mathcal{L}_{CL}$ . Refer to Figure 2 for the choice of hyperparameters.

Eq. 2 is a type of categorical cross-entropy loss, where the numerator (first and second terms) contains the positive pairs while the denominator (third term) consists of both positive and negative pairs. Our contrastive loss is clearly different to the conventional approach as it considers representations from mixed sources as positive pairs, rather than representations from multi-augmented images [25, 26]. Through the mixing ratio  $\lambda$ , the model learns the degree to which mixed representations should contain the inherent information of the original samples.

As summarized in Table 2, our experimental findings demonstrate that contrasting Patch-Mixed representations among all batch samples, while applying stop-gradient to the target representations (e.g.,  $z_i$  or  $z_m$ ), yielded the most promising outcomes. In practice, our proposed Patch-Mix Contrastive Learning significantly improves the Score on the ICBHI dataset by 2.82% compared to the simple fine-tuning of the AST. Given that the conventional mixed loss function did not improve the performance, it proves that our mixed contrastive learning in the latent space is highly effective.

## 5. ICBHI Dataset Results

Table 3 presents the overall respiratory sound classification performances of the ICBHI dataset, including our proposed method. Based on the overall comparisons, pretrained models on large-scale datasets showed relatively high performance regardless of architectures or methods. Interestingly, the straight-

forward fine-tuning of the AST model (“Fine-tuning”) achieves a 59.55% Score, which outperforms the current state-of-the-art model [9] by 1.26%. In addition, our Patch-Mix Contrastive Learning (“Patch-Mix CL”) achieved the state-of-the-art Score in both 4-class and 2-class evaluations, with the values of 62.37% and 68.71%, respectively. We would like to highlight that our methodology demonstrated a significant increase of approximately 3% compared to simple fine-tuning of the AST, while the methodologies proposed in previous literature have shown only a 1-2% increase within their architectures.

## 6. Conclusion

In this work, we demonstrate the effectiveness of the AST model, a purely attention-based model pretrained on visual and audio datasets, in improving the classification Score of respiratory sounds. Additionally, we propose a Patch-Mix augmentation that randomly replaces patches with those from different samples. However, the conventional mixed loss function did not result in performance gain, likely due to the label hierarchy in lung sounds. Therefore, we introduce a Patch-Mix Contrastive Learning that distinguishes the mixed representations in the latent space. Our method achieved a new record 62.37% Score on the ICBHI dataset, outperforming all existing methods.

## 7. Acknowledgements

This work was supported by the Ministry of Employment and Labor and HRD Korea’s K-Digital Platform project, and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2019-0-00075, Artificial Intelligence Graduate School Program (KAIST) and No.2022-0-00641, XVoice: Multi-Modal Voice Meta Learning), and by the MSIT (Ministry of Science and ICT), Korea, under the ITRC (Information Technology Research Center) support program (IITP-2023-2020-0-01808) supervised by the IITP.

## 8. References

- [1] Forum of International Respiratory Societies, “The global impact of respiratory disease. third edition.” 2021. [Online]. Available: [firsnet.org/images/publications/FIRS\\_Master\\_09202021.pdf](https://firsnet.org/images/publications/FIRS_Master_09202021.pdf)
- [2] Y. M. Arabi, E. Azoulay, H. M. Al-Dorzi, J. Phua, J. Salluh, A. Binnie, C. Hodgson, D. C. Angus, M. Cecconi, B. Du *et al.*, “How the covid-19 pandemic will change the future of critical care,” *Intensive care medicine*, vol. 47, pp. 282–291, 2021.
- [3] M. Sarkar, I. Madabhavi, N. Niranjana, and M. Dogra, “Auscultation of the respiratory system,” *Annals of thoracic medicine*, vol. 10, no. 3, p. 158, 2015.
- [4] L. Arts, E. H. T. Lim, P. M. van de Ven, L. Heunks, and P. R. Tuinman, “The diagnostic accuracy of lung auscultation in adult patients with acute pulmonary pathologies: a meta-analysis,” *Scientific reports*, vol. 10, no. 1, pp. 1–11, 2020.
- [5] B. Rocha, D. Filos, L. Mendes, I. Vogiatzis, E. Perantoni, E. Kaimakamis, P. Natsiavas, A. Oliveira, C. Jácome, A. Marques *et al.*, “A respiratory sound database for the development of automated classification,” in *Precision Medicine Powered by pHealth and Connected Health: ICBHI 2017, Thessaloniki, Greece, 18-21 November 2017*. Springer, 2018, pp. 33–37.
- [6] Y. Ma, X. Xu, and Y. Li, “Lungrn+ nl: An improved adventitious lung sound classification using non-local block resnet neural network with mixup data augmentation,” in *Interspeech*, 2020, pp. 2902–2906.
- [7] Z. Yang, S. Liu, M. Song, E. Parada-Cabaleiro, and B. W. Schuller, “Adventitious respiratory classification using attentive residual neural networks,” in *Interspeech*, 2020.
- [8] S. Gairola, F. Tom, N. Kwatra, and M. Jain, “Respirenet: A deep neural network for accurately detecting abnormal lung sounds in limited data setting,” in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2021, pp. 527–530.
- [9] T. Nguyen and F. Pernkopf, “Lung sound classification using co-tuning and stochastic normalization,” *IEEE Transactions on Biomedical Engineering*, vol. 69, no. 9, pp. 2872–2882, 2022.
- [10] A. Bohadana, G. Izbicki, and S. S. Kraman, “Fundamentals of lung auscultation,” *New England Journal of Medicine*, vol. 370, no. 8, pp. 744–751, 2014.
- [11] B. Flietstra, N. Markuzon, A. Vyshedskiy, and R. Murphy, “Automated analysis of crackles in patients with interstitial pulmonary fibrosis,” *Pulmonary medicine*, vol. 2011, 2011.
- [12] S. Reichert, R. Gass, C. Brandt, and E. Andrès, “Analysis of respiratory sounds: state of the art,” *Clinical medicine. Circulatory, respiratory and pulmonary medicine*, vol. 2, pp. CCRPM–S530, 2008.
- [13] Y. Ma, X. Xu, Q. Yu, Y. Zhang, Y. Li, J. Zhao, and G. Wang, “Lungbrn: A smart digital stethoscope for detecting respiratory disease using bi-resnet deep learning algorithm,” in *2019 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 2019, pp. 1–4.
- [14] Z. Wang and Z. Wang, “A domain transfer based data augmentation method for automated respiratory classification,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 9017–9021.
- [15] I. Moummed and N. Farrugia, “Supervised contrastive learning for respiratory sound classification,” *arXiv preprint arXiv:2210.16192*, 2022.
- [16] Y. Gong, Y.-A. Chung, and J. Glass, “AST: Audio Spectrogram Transformer,” in *Proc. Interspeech 2021*, 2021, pp. 571–575.
- [17] Y. Gong, C.-I. Lai, Y.-A. Chung, and J. Glass, “Ssast: Self-supervised audio spectrogram transformer,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 10, 2022, pp. 10 699–10 709.
- [18] A. Baade, P. Peng, and D. Harwath, “MAE-AST: masked autoencoding audio spectrogram transformer,” in *Interspeech 2022*, ISCA, 2022, pp. 2438–2442.
- [19] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [20] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *Proc. IEEE ICASSP 2017*, New Orleans, LA, 2017.
- [21] H. Zhang, M. Cissé, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” in *6th International Conference on Learning Representations, ICLR 2018*. OpenReview.net, 2018.
- [22] S. Yun, D. Han, S. J. Oh, S. Chun, J. Choe, and Y. Yoo, “Cutmix: Regularization strategy to train strong classifiers with localizable features,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 6023–6032.
- [23] S. Kim, G. Lee, S. Bae, and S.-Y. Yun, “Mixco: Mix-up contrastive learning for visual representation,” *arXiv preprint arXiv:2010.06300*, 2020.
- [24] K. Lee, Y. Zhu, K. Sohn, C. Li, J. Shin, and H. Lee, “i-mix: A domain-agnostic strategy for contrastive representation learning,” in *9th International Conference on Learning Representations, ICLR 2021*. OpenReview.net, 2021.
- [25] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [26] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, “Momentum contrast for unsupervised visual representation learning,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 9729–9738.
- [27] J. Li, J. Yuan, H. Wang, S. Liu, Q. Guo, Y. Ma, Y. Li, L. Zhao, and G. Wang, “Lungattn: advanced lung sound classification using attention mechanism with dual tqwt and triple stft spectrogram,” *Physiological Measurement*, vol. 42, no. 10, p. 105006, 2021.
- [28] J. Liu, B. Liu, H. Zhou, H. Li, and Y. Liu, “Tokenmix: Rethinking image mixing for data augmentation in vision transformers,” in *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVI*. Springer, 2022, pp. 455–471.
- [29] H. K. Choi, J. Choi, and H. J. Kim, “Tokenmixup: Efficient attention-guided token-level data augmentation for transformers,” in *Advances in Neural Information Processing Systems*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022.
- [30] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” *Interspeech 2019*, Sep 2019.
- [31] L. Xu, J. Cheng, J. Liu, H. Kuang, F. Wu, and J. Wang, “Arsnet: Adventitious respiratory sound classification network using parallel paths with channel-spatial attention,” in *2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, 2021, pp. 1125–1130.
- [32] L. Pham, H. Phan, R. Palaniappan, A. Mertins, and I. McLoughlin, “Cnn-moe based framework for classification of respiratory anomalies and lung disease detection,” *IEEE journal of biomedical and health informatics*, vol. 25, no. 8, pp. 2938–2947, 2021.
- [33] Z. Ren, T. T. Nguyen, and W. Nejdli, “Prototype learning for interpretable respiratory sound analysis,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 9087–9091.
- [34] Y. Chang, Z. Ren, T. T. Nguyen, W. Nejdli, and B. W. Schuller, “Example-based Explanations with Adversarial Attacks for Respiratory Sound Analysis,” in *Proc. Interspeech 2022*, 2022, pp. 4003–4007.
- [35] L. Pham, D. Ngo, K. Tran, T. Hoang, A. Schindler, and I. McLoughlin, “An ensemble of deep learning frameworks for predicting respiratory anomalies,” in *2022 44th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2022, pp. 4595–4598.